

Cover-Decomposition and Polychromatic Numbers*

Béla Bollobás[†] David Pritchard[‡] Thomas Rothvoß[§] Alex Scott[¶]

May 31, 2012

Abstract

A colouring of a hypergraph's vertices is *polychromatic* if every hyperedge contains at least one vertex of each colour; the *polychromatic number* is the maximum number of colours in such a colouring. Its dual, the *cover-decomposition number*, is the maximum number of disjoint hyperedge-covers. In geometric hypergraphs, there is extensive work on lower-bounding these numbers in terms of their trivial upper bounds (minimum hyperedge size and degree); our goal here is to broaden the study beyond geometric settings. We obtain algorithms yielding near-tight bounds for three families of hypergraphs: bounded hyperedge size, paths in trees, and bounded VC-dimension. This reveals that discrepancy theory and iterated linear program relaxation are useful for cover-decomposition. Finally, we discuss the generalization of cover-decomposition to sensor cover.

1 Introduction

In a set system on vertex set V , a subsystem is a *set cover* if each vertex of V appears in at least 1 set of the subsystem. Suppose that in the whole system, each vertex appears in at least δ sets, for some large δ ; does it follow that we can partition the system into 2 subsystems, such that each subsystem is a set cover?

Many natural families of set systems admit a universal constant δ for which this question has an affirmative answer. Such families are called *cover-decomposable*. But the family of *all* set systems is not cover-decomposable, as the following example shows. For any positive integer k , consider a set system which has $2k - 1$ sets, and where every k sets contain one mutually common vertex not contained by the other $k - 1$ sets. This system satisfies the hypothesis of the question for $\delta = k$. But every set cover has $\geq k$ sets, and since there are only $2k - 1$ sets in total, no partition into two set covers is possible. This example above shows that some sort of restriction on the family is necessary to ensure cover-decomposability.

One positive example of cover-decomposition arises if every set has size 2: such hypergraphs are simply graphs. They are cover-decomposable with $\delta = 3$: any graph with minimum degree 3 can have its edges partitioned into two edge covers. More generally, Gupta [20] showed (see also [1, 5]) that we can partition the edges of any multigraph into $\lfloor \frac{3\delta+1}{4} \rfloor$ edge covers. This bound is tight, even for 3-vertex multigraphs. (In *simple* graphs the optimal bound [19] is $\delta - 1$.)

Set systems in many geometric settings have been studied with respect to cover-decomposability; many positive and negative examples are known and there is no easy way to distinguish one from the other. In the affirmative case, as with Gupta's theorem, the next natural problem is to find for each $t \geq 2$ the smallest $\delta(t)$ such that when each vertex appears in at least $\delta(t)$ sets, a partition into t set covers is possible. The goal of this paper is to extend the study of cover-decomposition beyond geometric settings.

*A preliminary version of this paper appeared as [9].

[†]University of Memphis, USA and University of Cambridge, UK; supported in part by NSF grants CNS-0721983, CCF-0728928, DMS-0906634 and CCR-0225610, and ARO grant W911NF-06-1-0076.

[‡]University of Waterloo, Canada. Work performed while at École Polytechnique Fédérale de Lausanne, Switzerland and supported by an NSERC Post-Doctoral Fellowship.

[§]MIT, Cambridge, USA; supported by the Alexander von Humboldt Foundation within the Feodor Lynen program, by ONR grant N00014-11-1-0053, and by NSF contract CCF-0829878.

[¶]Mathematical Institute, University of Oxford, UK.

1.1 Terminology and Notation

A *set system* or *hypergraph* $H = (V, \mathcal{E})$ consists of a ground set V of vertices, together with a collection \mathcal{E} of hyperedges, where each hyperedge $E \in \mathcal{E}$ is a subset of V . We will sometimes call hyperedges just *edges* or *sets*. We permit \mathcal{E} to contain multiple copies of the same hyperedge (e.g. to allow us to define “duals” and “shrinking” later), and we also allow hyperedges of cardinality 0 or 1. We only consider hypergraphs that are finite. Note, in some geometric cases, infinite cover-decomposability problems can be reduced to finite ones. We refer the reader to [34] including the distinction between plane- and total-cover-decomposability. Additional work on the infinite version appears in [14].

To *shrink* a hyperedge E in a hypergraph means to replace it with some $E' \subseteq E$. This operation is useful in several places.

A *polychromatic k -colouring* of a hypergraph is a function from V to a set of k colours so that for every edge, its image contains all colours. Equivalently, the colour classes partition V into sets which each meet every edge, so-called *vertex covers/transversals*. The maximum number of colours in a polychromatic colouring of H is called its *polychromatic number*, which we denote by $\mathbf{p}(H)$.

A *cover k -decomposition* of a hypergraph is a partition of \mathcal{E} into k subfamilies $\mathcal{E} = \bigsqcup_{i=1}^k \{\mathcal{E}_i\}$ such that each $\bigcup_{E \in \mathcal{E}_i} E = V$. In other words, each \mathcal{E}_i must be a set cover. The maximum k for which the hypergraph H admits a cover k -decomposition is called its *cover-decomposition number*, which we denote by $\mathbf{cd}(H)$.

The *dual* H^* of a hypergraph H is another hypergraph such that the vertex set of H^* corresponds to the edge set of H , and vice-versa, with incidences preserved. Thus the vertex-edge incidence matrices for H and H^* are transposes of one another. E.g., the standard notation for the example in the introduction is $\binom{[2k-1]}{k}^*$. From the definitions it is easy to see that the polychromatic and cover-decomposition numbers are dual to one another,

$$\mathbf{cd}(H) = \mathbf{p}(H^*).$$

The *degree* of a vertex v in a hypergraph is the number of hyperedges containing v ; it is *d -regular* if all vertices have degree d . We denote the minimum degree by δ , and the maximum degree by Δ . We denote the minimum size of any hyperedge by r , and the maximum size of any hyperedge by R . Note that $\Delta(H) = R(H^*)$ and $\delta(H) = r(H^*)$. It is trivial to see that $\mathbf{p} \leq r$ in any hypergraph and dually that $\mathbf{cd} \leq \delta$. So the cover-decomposability question asks if there is a converse to this trivial bound: if δ is large enough, does \mathbf{cd} also grow? To write this concisely, for a family \mathcal{F} of hypergraphs, let its extremal cover-decomposition function $\overline{\mathbf{cd}}(\mathcal{F}, \delta)$ be

$$\overline{\mathbf{cd}}(\mathcal{F}, \delta) := \min\{\mathbf{cd}(H) \mid H \in \mathcal{F}; \forall v \in V(H) : \text{degree}(v) \geq \delta\},$$

i.e. $\overline{\mathbf{cd}}(\mathcal{F}, \delta)$ is the best possible lower bound for \mathbf{cd} among hypergraphs in \mathcal{F} with min-degree $\geq \delta$. So to say that \mathcal{F} is cover-decomposable means that $\overline{\mathbf{cd}}(\mathcal{F}, \delta) > 1$ for some constant δ . We also dually define

$$\overline{\mathbf{p}}(\mathcal{F}, r) := \min\{\mathbf{p}(H) \mid H \in \mathcal{F}; \forall E \in \mathcal{E}(H) : |E| \geq r\}.$$

In the rest of the paper we focus on computing the functions $\overline{\mathbf{cd}}$ and $\overline{\mathbf{p}}$.

We sometimes write \mathbf{p} for $\mathbf{p}(H)$ when H is clear from context, and likewise with $\overline{\mathbf{p}}$. We write $\overline{\mathbf{cd}}(\delta)$ or $\overline{\mathbf{p}}(r)$ when the family \mathcal{F} is clear from context. For example, Gupta’s theorem (together with the tight example) says that in graphs, $\overline{\mathbf{cd}}(\delta) = \lfloor \frac{3\delta+1}{4} \rfloor$.

1.2 Results

In Section 2 we generalize Gupta’s theorem to hypergraphs of bounded edge size. Let $\text{HYP}(R)$ denote the family of hypergraphs with all edges of size at most R .

Theorem 1. *For all R, δ we have $\overline{\mathbf{cd}}(\text{HYP}(R), \delta) \geq \delta/(\ln R + O(\ln \ln R))$.*

Before proving Theorem 1, we give a simpler proof of a result that is weaker by a constant factor. This proof uses the Lovász Local Lemma (LLL) and the Beck-Fiala theorem [7]. The Beck-Fiala theorem says

that every hypergraph's *discrepancy* (defined later) is less than 2Δ . To get the strong version of Theorem 1 we employ another discrepancy upper bound of $2\sqrt{R \ln(R\Delta)}$, obtained using the LLL and Chernoff bounds (Proposition 12). Next we show that Theorem 1 is always tight up to a constant factor, and that when R is sub-exponential in δ this constant factor tends to 1.

Theorem 2. (i) For all $R \geq 2, \delta \geq 1$ we have $\overline{\text{cd}}(\text{HYP}(R), \delta) \leq \max\{1, O(\delta/\ln R)\}$.

(ii) For any sequence $R, \delta \rightarrow \infty$ with $\delta = \omega(\ln R)$ we have $\overline{\text{cd}}(\text{HYP}(R), \delta) \leq (1 + o(1))\delta/\ln(R)$.

Here (i) uses an explicit construction while (ii) uses the probabilistic method. By plugging Theorem 1 into an approach of [1], one obtains a good bound on the cover-decomposition number of *sparse* hypergraphs.

Corollary 3. Suppose for some fixed α, β that $H = (V, \mathcal{E})$ satisfies, for all $V' \subseteq V$ and $\mathcal{E}' \subseteq \mathcal{E}$, that the number of incidences between V' and \mathcal{E}' is at most $\alpha|V'| + \beta|\mathcal{E}'|$. Then $\text{cd}(H) \geq \frac{\delta(H) - \alpha}{\ln \beta + O(\ln \ln \beta)}$.

Note that duality yields a similar bound on the polychromatic number.

In Section 3 we consider the following family of hypergraphs: the ground set is the edge set of an undirected tree, and each hyperedge must correspond to the edges lying in some path in the tree. We show that such systems are cover-decomposable:

Theorem 4. For hypergraphs defined by edges of paths in trees, $\overline{\text{cd}}(\delta) = \Omega(\delta)$.

To prove Theorem 4 we exploit iterated LP relaxation, using an extreme point structure theorem for paths in trees from [24]. The approach also generalizes to path-in-tree systems where the ground set consists of the arcs or the nodes of the tree. We also determine the extremal polychromatic number for such systems:

Theorem 5. For hypergraphs defined by edges of paths in trees, $\overline{\text{p}}(r) = \lceil r/2 \rceil$.

This contrasts with a construction of Pach, Tardos and Tóth [32] (described in Section 4): if we also allow hyperedges consisting of sets of “siblings,” then $\overline{\text{p}}(r) = 1$ for all r .

The *VC-dimension* is a prominent measure of set system complexity used frequently in geometry: it is the maximum cardinality of any $S \subseteq V$ such that $\{S \cap E \mid E \in \mathcal{E}\} = 2^S$. It is natural to ask what role, if any, the VC-dimension plays in cover-decomposability. In the most restricted case, we get:

Theorem 6. For the family of hypergraphs with VC-dimension 1, $\overline{\text{p}}(r) = \lceil r/2 \rceil$ and $\overline{\text{cd}}(\delta) = \lceil \delta/2 \rceil$.

By duality, the same holds for the family of hypergraphs whose duals have VC-dimension 1. Roughly, these results hold because VC-dimension 1 implies a strong structural property. Moreover, we observe that a dimension upper bound of 2 does not yield cover-decomposability:

Theorem 7. For the family of hypergraphs $\{H \mid \text{VC-dim}(H), \text{VC-dim}(H^*) \leq 2\}$, we have $\overline{\text{p}}(r) = 1$ for all r and $\overline{\text{cd}}(\delta) = 1$ for all δ .

To prove this, we show that the construction of [32] has primal and dual VC-dimension at most 2.

In the last section, we discuss the generalization of cover-decomposition to cover-scheduling. This is inspired by results from [18], where a proof of $\text{cd} = \Omega(\delta)$ for convex polygons in the plane is generalized to a scheduling setting while losing only a constant factor. We do not know if Theorem 1 has an analogue within a constant factor; we prove a weaker version with linear dependence on R .

All of our lower bounds on $\overline{\text{p}}$ and $\overline{\text{cd}}$ can be implemented as polynomial-time algorithms. In the case of Theorem 1 this relies on the constructive LLL framework of Moser-Tardos [28]. In the tree setting (Theorem 4) this relies on a linear programming subroutine. Note: since we also have the trivial bounds $\text{p} \leq r$ and $\text{cd} \leq \delta$, these give approximation algorithms for p and cd , e.g. Theorem 1 gives a $(\ln R + O(\ln \ln R))$ -approximation for cd .

1.3 Related Work

One practical motive to study cover-decomposition is that the hypergraph can model a collection of sensors [10, 18], with each $E \in \mathcal{E}$ corresponding to a sensor which can monitor the subset $E \subset V$ of vertices; then monitoring all of V takes a set cover, and $\overline{\text{cd}}$ is the maximum “coverage” of V possible if each sensor can only be turned on for a single time unit or monitor a single frequency. A motivation from theory is that if $\overline{\text{cd}}(\delta) = \Omega(\delta)$ holds for a family closed under vertex deletion, then the size of a *dual ϵ -net* is bounded by $O(1/\epsilon)$ [31].

A hypergraph is said to be *weakly k -colourable* if we can k -colour its vertex set so that no edge is monochromatic. Weak 2-colourability is also known as *Property B*, and these notions coincide with the property $\mathbf{p} \geq 2$. However, weak k -colourability does not imply $\mathbf{p} \geq k$ in general. The *discrepancy* of a hypergraph is the minimum D so that we can colour its vertices by ± 1 so that the sum of the values within every edge is between $-D$ and D . This is related to Property B: if a hypergraph’s discrepancy is strictly less than δ , then the same colouring establishes that it has Property B. We will deal with discrepancy in the dual setting.

For a graph $G = (V, E)$, the (*closed*) *neighbourhood hypergraph* $\mathcal{N}(H)$ is defined to be a hypergraph on ground set V , with one hyperedge $\{v\} \cup \{u \mid \{u, v\} \in E\}$ for each $v \in V$. Then $\mathbf{p}(\mathcal{N}(G)) = \text{cd}(\mathcal{N}(G))$ equals the *domatic number* of G , i.e. the maximum number of disjoint dominating sets. The paper of Feige, Halldórsson, Kortsarz & Srinivasan [16] obtains upper bounds for the domatic number and their bounds are essentially the same as what we get by applying Theorem 1 to the special case of neighbourhood hypergraphs; compared to our methods they use the LLL but not discrepancy or iterated LP relaxation. They give a hardness-of-approximation result which implies that Theorem 1 is tight with respect to the approximation factor, namely for all $\epsilon > 0$, it is hard to approximate cd within a factor better than $(1 - \epsilon) \ln R$, under reasonable complexity assumptions. A simpler, weaker hardness result is that you cannot approximate cd of graphs better than $3/2$, since it is NP-hard to check if a cubic graph has 3 disjoint perfect matchings [21]. A generalization of results in [16] to packing polymatroid bases was given in [12]; this implies a weak version of Theorem 1 where the $\ln R$ term is replaced by $\ln |V|$.

Given a plane graph, define a hypergraph whose vertices are the graph’s vertices, and whose hyperedges are its faces. In [1] it was shown to satisfy the hypothesis of Corollary 3 for $\alpha = \beta = 2$. Their method, which we re-use in the proof of Corollary 3, gives $\overline{\mathbf{p}}(\delta) \leq \lfloor (3\delta - 5)/4 \rfloor$ from Gupta’s theorem.

A notable progenitor in geometric literature on cover-decomposition is the following question of Pach [29]. Take a convex set $A \subset \mathbb{R}^2$. Let $\mathbb{R}^2|\text{TRANSLATES}(A)$ denote the family of hypergraphs where the ground set V is a finite subset of \mathbb{R}^2 , and each hyperedge is the intersection of V with some translate of A . Pach asked if such systems are cover-decomposable, and this question is still open. A state-of-the-art partial answer is due to Gibson & Varadarajan [18], who prove that $\overline{\mathbf{p}}(\mathbb{R}^2|\text{TRANSLATES}(A), \delta) = \Omega(\delta)$ when A is an open convex polygon; prior work includes [30, 40, 33, 4, 36]. There is an unpublished proof [26] that $\overline{\mathbf{p}}(\mathbb{R}^2|\text{UNIT-DISCS}, 33) \geq 2$. On the other hand, unit balls in \mathbb{R}^3 or higher dimensions are not cover-decomposable [27, 32].

Pach, Tardos and Tóth [32] obtain several negative results by embedding a non-cover-decomposable tree-based construction (used in Section 4) in geometric settings. In this way they prove that the following families are not cover-decomposable: $\mathbb{R}^2|\text{AXIS-ALIGNED-RECTANGLES}$; $\mathbb{R}^2|\text{TRANSLATES}(A)$ when A is a quadrilateral that is not convex; and $\mathbb{R}^2|\text{STRIPS}$ and its dual. In contrast to the latter result, it is known that $\overline{\mathbf{p}}(\mathbb{R}^2|\text{AXIS-ALIGNED-STRIPS}, r) \geq \lceil r/2 \rceil$ [3]. Recently it was shown [23] that $\mathbb{R}^3|\text{TRANSLATES}(\mathbb{R}_+^3)$ is cover-decomposable, giving cover-decomposability of $\mathbb{R}^2|\text{HOMOTHESTS}(T)$ for any triangle T and a new proof (c.f. [22]) for $\mathbb{R}^2|\text{BOTTOMLESS-AXIS-ALIGNED-RECTANGLES}$; the former contrasts with the fact that $\mathbb{R}^2|\text{DISCS}$ is not cover-decomposable [32].

For $\mathbb{R}^2|\text{HALFSPACES}$ several results are known. Smorodinsky and Yuditsky [39] proved $\overline{\text{cd}}(\delta) = \lceil \delta/2 \rceil$ (improving upon [6]) and $\overline{\mathbf{p}}(r) \geq \lceil r/3 \rceil$. They note the latter is not tight, since Fulek [17] showed $\overline{\mathbf{p}}(3) = 2$.

Aside from the tree-based construction of [32], we mention two other indecomposable constructions. First, the Hales-Jewett theorem is used in [32] to show that $\overline{\mathbf{p}}(\mathbb{R}^2|\text{LINES}, r) = 1$ for all r . Second, [35] gives the following indecomposable construction, which is smaller than that of [32]. The ground set of the hypergraph is the set of all strings on at most $r - 1$ B’s and at most $r - 1$ R’s. For every c in $\mathbf{Z}_{\geq 0}^r$ with $\sum_i c[i] \leq r - 1$, there

is a hyperedge $\{R^{c[1]}, R^{c[1]}BR^{c[2]}, \dots, R^{c[1]}BR^{c[2]}B \dots BR^{c[r]}\}$, and another hyperedge obtained by swapping the roles of B and R. Then [35] shows such hypergraphs have $p = 1$.

A recent paper of Chan et al. [11] centers around *quasi-uniform* distributions over set covers, which are fractional packings of set covers. In contrast, cd deals with integral packing of set covers. The inspiration in [11] is rooted in geometric settings, but they arrive at general combinatorial results.

For a hypergraph H let M_H be its 0-1 incidence matrix, with rows for edges and columns for vertices. Then the following three properties are equivalent [38]: (i) for every hypergraph obtained from H by deleting or duplicating vertices (columns), $p = r$; (ii) the fractional vertex cover polytope $\{x \mid x \geq 0, M_H x \geq 1\}$ has the *integer decomposition property*; (iii) the *blocker* of H is *Mengerian*. A special case of such hypergraphs are *balanced* hypergraphs, where no submatrix of M_H is the incidence matrix of an odd cycle.

A generalization of polychromatic c -colouring to c -strong colouring is considered in [8]: every hyperedge of size at least c must receive all c colours, and every smaller hyperedge must receive no colour twice. This can be further generalized to requiring an arbitrary lower bound $f(S)$ on the number of colours for every set S ; when f is the maximum of two supermodular functions this is the classical *supermodular colouring* problem solved originally by Schrijver [37].

2 Hypergraphs of Bounded Edge Size

To get good upper bounds on $\overline{\text{cd}}(\text{HYP}(R), \delta)$, we will use the Lovász Local Lemma (LLL):

Lemma 8 (LLL, [15]). *Consider a collection of “bad” events such that each one has probability at most p , and such that each bad event is independent of the other bad events except at most D of them. (We call D the dependence degree.) If $p(D+1)e \leq 1$ then with positive probability, no bad events occur.*

Our first tool extends a standard argument about Property B [2, Theorem 5.2.1].

Proposition 9. $\overline{\text{cd}}(\text{HYP}(R), \delta) \geq \lfloor \delta / \ln(eR\delta^2) \rfloor$.

Proof. Given any hypergraph $H = (V, \mathcal{E})$ where every edge has size at most R and such that each $v \in V$ is covered at least δ times, we must show for $t = \lfloor \delta / \ln(eR\delta^2) \rfloor$ that $\text{cd}(H) \geq t$, i.e. that \mathcal{E} can be decomposed into t disjoint set covers. It will be helpful here and later to make the degree of every vertex *exactly* δ , since this bounds the dependence degree. (A similar observation was made by Mani-Levitska and Pach [26] in the setting of cover-decomposing discs.) This is without loss of generality: otherwise whenever $\deg(v) > \delta$ shrink some $E \ni v$ to $E \setminus \{v\}$ until $\deg(v)$ drops to δ ; then observe that undoing shrinking preserves the property of being a set cover.

Consider the following randomized experiment: for each hyperedge $E \in \mathcal{E}$, assign a random colour between 1 and t to E . If we can show that with positive probability, every vertex is incident with a hyperedge of each colour, then we will be done.

For each vertex v define the *bad event* \mathfrak{E}_v to be the event that v is not incident with a hyperedge of each colour. The probability of \mathfrak{E}_v is at most $t(1 - \frac{1}{t})^\delta$, by using a union bound. The event \mathfrak{E}_v only depends on the colours of the hyperedges containing v ; therefore the events \mathfrak{E}_v and $\mathfrak{E}_{v'}$ are independent unless v, v' are in a common hyperedge. In particular the dependence degree is at most $(R-1)\delta < R\delta$, since each edge containing v contains at most $R-1$ other vertices. It follows by LLL that if

$$R\delta t(1 - \frac{1}{t})^\delta \leq 1/e,$$

then with positive probability, no bad events happen and we are done. We can verify that $t = \delta / \ln(eR\delta^2)$ satisfies this bound. \square

We will next show that the bound can be raised to $\Omega(\delta / \ln R)$. Intuitively, our strategy is the following. The lower bound on cd given by Proposition 9, $\delta / \ln(R\delta)$, is already $\Omega(\delta / \ln R)$ in the special case that $\delta \leq R^{O(1)}$. For hypergraphs where $\delta \gg R$ we will show that we can partition \mathcal{E} into m parts $\mathcal{E} = \bigsqcup_{i=1}^m \mathcal{E}_i$ so that $\delta(V, \mathcal{E}_i)$ is at least a constant fraction of δ/m , and such that δ/m is polynomial in R . Thus by

Proposition 9 we can extract $\Omega((\delta/m)/\ln R)$ set covers from each (V, \mathcal{E}_i) , and taken all together we obtain $\Omega(\delta/\ln R)$ set covers.

In fact, it will be enough to consider splitting \mathcal{E} into two parts at a time, recursively. Then ensuring $\delta(V, \mathcal{E}_i) \gtrsim \delta/2$ ($i = 1, 2$) amounts to a dual discrepancy-theoretic problem: we must 2-colour the hyperedges by ± 1 so that for each vertex, the sum of the incident hyperedges' colours is in $[-D, D]$, with the discrepancy D as small as possible. To get a short proof of a weaker version of Theorem 1, we use a theorem of Beck and Fiala [7]. Later, in Section 3, we will extend their approach to trees.

Theorem 10 (Beck & Fiala [7], stated in the dual). *In a δ -regular hypergraph $H = (V, \mathcal{E})$ with all edges of size at most R , we can partition the edge set into $\mathcal{E} = \mathcal{E}_1 \uplus \mathcal{E}_2$ such that $\delta(V, \mathcal{E}_i) \geq \delta/2 - R$ for each $i \in \{1, 2\}$.*

Here is how the Beck-Fiala Theorem gives a near-optimal bound on $\overline{\text{cd}}$.

Proposition 11. $\overline{\text{cd}}(\text{HYP}(R), \delta) \geq \Omega(\delta/\ln R)$.

Proof. If $\delta < 4R$ this already follows from Proposition 9. Otherwise apply Theorem 10 to the initial hypergraph, and then use shrinking to make both the resulting (V, \mathcal{E}_i) 's into regular hypergraphs. Iterate this process; stop splitting each hypergraph once its degree falls in the range $[R, 4R)$, which is possible since $\delta \geq 4R \Rightarrow \delta/2 - R \geq R$. Let M be the number of hypergraphs at the end.

Observe that in applying the splitting-and-shrinking operation to some (V, \mathcal{E}) to get (V, \mathcal{E}_1) and (V, \mathcal{E}_2) , the sum of the degrees of (V, \mathcal{E}_1) and (V, \mathcal{E}_2) is at least the degree of (V, \mathcal{E}) , minus $2R$ “waste”. It follows that the total waste is at most $2R(M-1)$, and we have that $4RM + 2R(M-1) \geq \delta$. Consequently $M \geq \delta/6R$. As sketched earlier, applying Proposition 9 to the individual hypergraphs, and combining these vertex covers, shows that $\text{cd} \geq M \lfloor R/\ln(eR^3) \rfloor$ which gives the claimed bound. \square

Now we get to the better bound with the correct multiplicative constant. We will need the following proposition, which when stated in the dual setting, says that the discrepancy of a hypergraph is at most $2\sqrt{R \ln(R\Delta)}$.

Proposition 12. *In a hypergraph $H = (V, \mathcal{E})$ with maximum degree Δ and maximum edge size R , we can partition the edge set into $\mathcal{E} = \mathcal{E}_1 \uplus \mathcal{E}_2$ such that for every vertex v , the degree of v in each \mathcal{E}_i is at least $\text{degree}_H(v)/2 - \sqrt{\Delta \ln(R\Delta)}$, assuming $R\Delta > 1$.*

Proof. Let d_v be short for $\text{degree}_H(v)/2$. Independently assign each edge to \mathcal{E}_1 or \mathcal{E}_2 uniformly at random. For a $\lambda \geq 0$ to be fixed later, for each vertex, let the event \mathfrak{E}_v be that the degree of v is less than $\text{degree}_H(v)/2 - \lambda$ in some \mathcal{E}_i . Chernoff bounds [2, Corollary A.1.2] imply that \mathfrak{E}_v has probability at most $2\exp(-2\lambda^2/d_v) \leq 2\exp(-2\lambda^2/\Delta)$.

If with positive probability no bad events happen, then we are done. As in the proof of Proposition 9, the dependence degree is strictly less than $R\Delta$. Therefore the LLL is applicable provided that $2\exp(-2\lambda^2/\Delta)R\Delta \leq 1/e$, which is satisfied by $\lambda = \sqrt{\Delta \ln(2eR\Delta)}/2$. This gives the desired result under the condition that $R\Delta \geq 2e$. To finish the proof, one must check the cases $R = 1 < \Delta$, $\Delta = 1 < R$, and $\Delta = R = 2$, which are straightforward. \square

Now we get to the main proof.

Proof of Theorem 1. We want to show that for all hypergraphs with minimum degree δ and maximum hyperedge size R , that $\text{p} \geq \delta/(\ln R + O(\ln \ln R))$. Due to the crude bound in Proposition 11, we may assume that R is larger than any chosen fixed constant. Moreover, Proposition 9 gives us the desired bound when δ is at most polylogarithmic in R , so we assume $\delta \geq \ln^K R$ for a positive constant K to be fixed later. By shrinking, we assume the hypergraph is δ -regular.

Let $d_0 = \delta$ and $d_{i+1} = d_i/2 - \sqrt{d_i \ln(Rd_i)}$. We mimic the split-and-shrink proof of Proposition 11, using Proposition 12 for the splitting. After i rounds, all 2^i hypergraphs are regular with degree at least d_i . We stop splitting after T rounds, where T will be fixed later to make d_T and $\delta/2^T$ polylogarithmic in R . Using

the evident bound $d_i \leq \delta/2^i$, the total degree loss due to splitting is at most

$$\begin{aligned} \delta - 2^T d_T &= \sum_{i=0}^{T-1} 2^i (d_i - 2d_{i+1}) = \sum_{i=0}^{T-1} 2^i 2 \sqrt{d_i \ln R d_i} \\ &\leq \sum_{i=0}^{T-1} 2^i 2 \sqrt{(\delta/2^i) \ln(R\delta/2^i)} = 2\sqrt{\delta} \sum_{i=0}^{T-1} \sqrt{2} \sqrt{\ln(R\delta/2^i)}. \end{aligned} \quad (1)$$

In this series, $\sqrt{2}^i$ is increasing in i , and $\sqrt{\ln(R\delta/2^i)}$ decreases with i but not too quickly; the ratio between the i th and $(i+1)$ th terms is

$$\sqrt{\ln(R\delta/2^i)} / \sqrt{\ln(R\delta/2^{i+1})} = \sqrt{1 + \ln 2 / \ln(R\delta/2^{i+1})} \leq \sqrt{1 + \ln_R 2} < 1.3$$

where the first inequality uses $\delta/2^{i+1} \geq 1$ and the second uses $R \geq 3$. Hence the sequence in (1) grows at least geometrically with ratio $\sqrt{2}/1.3$ and its sum is within a constant factor of the final $i = T-1$ term, using the geometric series sum formula.

We deduce $\delta - 2^T d_T = O(\sqrt{\delta 2^T \ln(R\delta/2^T)})$. Pick $K = 3$, and pick T such that $\delta/2^T$ is between $\ln^3 R$ and $\ln^3 R/2$, then we have

$$d_T \geq \delta/2^T \left(1 - O\left(\sqrt{2^T \delta^{-1} \ln(R\delta/2^T)}\right)\right) = \delta/2^T (1 - O(\ln^{-1}(R))).$$

Consequently with Proposition 9 we see that

$$\text{cd} \geq 2^T d_T / (\ln R + O(\ln \ln R)) \geq \delta (1 - O(\ln^{-1}(R))) / (\ln R + O(\ln \ln R))$$

which gives the claimed bound.

2.1 Sparse Hypergraphs: Proof of Corollary 3

We re-use the following fact, a consequence of Hall's theorem, that was used in [1]; see also [25, Thm. 2.4.2].

Fact 13. *A hypergraph has at most $\alpha|V'| + \beta|\mathcal{E}'|$ incidences between each $V' \subseteq V$ and $\mathcal{E}' \subseteq \mathcal{E}$, if and only if the following assignment problem has a valid solution: for every e and every $v \in e$, we need to assign the incidence (v, e) to either v or e , and the number of total incidences assigned to each vertex (resp. edge) is at most α (resp. β).*

To prove Corollary 3, we apply the fact, and then use shrinking to remove the incidences assigned to vertices. We retain minimum degree $\delta - \alpha$ and maximum edge size β . Finally, the corollary follows from Theorem 1.

2.2 Lower Bounds

Now we show that the bounds obtained previously are tight.

Proof of Theorem 2(i). We want to show that for some constant C , all $R \geq 2$, and all $\delta \geq 1$, we have $\overline{\text{cd}}(\text{HYP}(R), \delta) \leq \max\{1, C\delta/\ln R\}$. Since $\overline{\text{cd}}(\text{HYP}(R), \delta)$ is non-increasing in R , we may reduce R by a constant factor to assume that $R = \binom{2k-2}{k-1}$ for some integer $k \geq 2$. Note this gives $k = \Theta(\ln R)$.

Consider the hypergraph $H = \binom{[2k-1]}{k}^*$ in the introduction. It is k -regular, it has $\text{cd}(H) = 1$, and $R(H) = \binom{2k-2}{k-1} = R$. If $\delta \leq k$ then H proves the theorem for small enough C , so assume $\delta \geq k$. Again by monotonicity, we may increase δ by a constant factor to make δ a multiple of k . Let $\mu = \delta/k$.

Consider the hypergraph μH obtained by copying each of its edges μ times, for an integer $\mu \geq 1$; note that it is δ -regular. The argument in the introduction shows that any set cover has size at least k and therefore average degree at least $k \binom{2k-2}{k-1} / \binom{2k-1}{k} = k^2 / (2k-1) = \Theta(\ln R)$. Thus $\overline{\text{cd}}(\mu H) = O(\delta/\ln R)$ which proves the theorem.

Next, Theorem 2(ii) establishes the right multiplicative constant in the range $\delta = \omega(\ln R)$. Our approach relies on the methods for neighbourhood hypergraphs established in [16]; in fact using [16] as a black box gives Theorem 2(ii) for the case $R \sim \delta$. As part of reaching the fuller range, Claim 14 will generalize a calculation [16, §2.5] for dominating sets in Erdős-Rényi graphs.

Proof of Theorem 2(ii). We must show, given a sequence $\{(R_i, \delta_i)\}_i$ such that $R \rightarrow \infty, \delta \rightarrow \infty$, and $\delta = \omega(\ln R)$ as $i \rightarrow \infty$, we have $\overline{\text{cd}}(\text{HYP}(R), \delta) \leq (1 + o(1))\delta / \ln(R)$. We assume an additional hypothesis, that $R \geq \delta$; this will be without loss of generality as we can handle the case $\delta > R$ using the μ -replication trick from the proof of Theorem 2(i), since our argument is again based on lower-bounding the minimum size of a set cover.

Let $\delta' = \delta(1 + o(1))$ and $R' = R(1 - o(1))$ be parameters that will be specified shortly. We construct a random hypergraph with $n = R'^2 \delta'$ vertices and $m = R' \delta'^2$ edges, where for each vertex v and each edge E , we have $v \in E$ with independent probability $p = 1/R' \delta'$. Thus each vertex has expected degree δ' and each edge has expected size R' . A standard Chernoff bound together with $np = \omega(\ln m)$ shows the maximum edge size is $(1 + o(1))R'$ asymptotically almost surely (a.a.s.); pick R' such that this $(1 + o(1))R'$ equals R . Likewise, since $mp = \omega(\ln n)$ we may pick $\delta' = (1 + o(1))\delta$ so that a.a.s. the minimum degree is at least δ . Now we prove that small set covers are not likely to exist:

Claim 14. *A.a.s. the minimum set cover size is at least $\frac{1}{p} \ln(pn)(1 - o(1))$.*

Proof. Fix $\epsilon > 0$ and let us a.a.s. lower-bound the minimum set cover size by $\frac{1}{p} \ln(pn)(1 - \epsilon)$. For any fixed collection of $s = \frac{1}{p} \ln(pn)(1 - \epsilon)$ edges, let P be the probability that it is a set cover. In the limit, p tends to 0, so $(1 - p) = \exp(-p + \Theta(p^2))$ and

$$\begin{aligned} P &= (1 - (1 - p)^s)^n = (1 - \exp(-s(p + \Theta(p^2))))^n \\ &= (1 - \exp(-(1 + \Theta(p))(1 - \epsilon) \ln(pn))))^n \\ &= (1 - (pn)^{-1 + \epsilon - \Theta(p)})^n. \end{aligned}$$

Since $1 - x \leq \exp(-x)$ for all x ,

$$P \leq \exp(-n(pn)^{-1 + \epsilon - \Theta(p)}) = \exp(-p^{-1}(pn)^{\epsilon - \Theta(p)}).$$

There are $\binom{m}{s} \leq m^s = \exp(s \ln m)$ such collections, so a.a.s. none are set covers provided that

$$s \ln m - p^{-1}(pn)^{\epsilon - \Theta(p)} \rightarrow -\infty.$$

In turn, expanding the definition of s , we need precisely that

$$p^{-1}(\ln(pn)(1 - \epsilon) \ln m - (pn)^{\epsilon - \Theta(p)}) \rightarrow -\infty.$$

To see this, note that $p = o(1)$, $pn = \omega(1)$, and m is polynomial in pn (since $R \geq \delta$). □

Using Claim 14 we can complete the proof of Theorem 2(ii), since the former implies that the maximum number of disjoint set covers cd is at most $(1 + o(1))mp / \ln(pn) = (1 + o(1))\delta' / \ln(R') = (1 + o(1))\delta / \ln(R)$.

Aside from the results above, not much else is known about specific values of $\overline{\text{cd}}(\text{HYP}(R), \delta)$ for small R, δ . The Fano plane gives $\overline{\text{p}}(\text{HYP}(3), 3) = 1$: if its seven sets are partitioned into two parts, one part has only three sets, and it is not hard to verify the only covers consisting of three sets are pencils through a point and therefore preclude the remaining sets from forming a cover. Moreover, Thomassen [41] showed that every 4-regular, 4-uniform hypergraph has Property B; together with monotonicity we deduce that $\overline{\text{p}}(\text{HYP}(3), 4) \geq \overline{\text{p}}(\text{HYP}(4), 4) \geq 2$.

3 Paths in Trees

Let TREEEDGES|PATHS denote the family of all hypergraphs that can be expressed in the following form: the ground set is the edge set E_T of some undirected tree T , and for each hyperedge E of H , there are vertices $u, v \in V_T$ so that E is the set of all edges in T 's u -to- v path. Note that interval hypergraphs correspond to the special case that the tree is itself a path graph.

Our main result for these tree hypergraphs is that they are cover-decomposable, and moreover the cover-decomposition number is linear in δ . To prove this theorem we use a counting lemma [24, Lemma 4] for linear program extreme points. While [24, Lemma 4] is stated for a linear program with just upper bounds on the capacity of edges, what we need is a result when there are both upper and lower bounds. However, the proof method is identical, as we will explain.

We give a simpler sketch of the main idea just for intuition. We will write a linear program about two-colouring paths so as to achieve small dual discrepancy. For every path/hyperedge P let x_P and y_P be non-negative variables constrained by $x_P + y_P = 1$. The idea is that $(x_P, y_P) = (1, 0)$ corresponds to placing P in \mathcal{E}_1 , and $(x_P, y_P) = (0, 1)$ corresponds to placing P in \mathcal{E}_2 . Consider the linear program

$$0 \leq x, y; x + y = 1; \forall e, \sum_{P: e \in P} x_P = \text{degree}_H(e)/2; \forall e, \sum_{P: e \in P} y_P = \text{degree}_H(e)/2.$$

If we could find an feasible integral solution (x, y) , then every vertex would be covered an equal number of times by \mathcal{E}_1 and \mathcal{E}_2 , and the dual discrepancy would be 0. There is always a feasible fractional solution $x = y = \frac{1}{2}$, but it is not guaranteed that a feasible integral solution exists. But, we can prove that there is an integer solution with $\pm O(1)$ additive violation in each constraint. In the real proof, we will eliminate the y variables according to $y = 1 - x$, and seek an unequal split in order to get a better factor in the final result.

Theorem 4. $\overline{\text{cd}}(\text{TREEEDGES|PATHS}, \delta) \geq 1 + \lfloor (\delta - 1)/5 \rfloor$.

Proof. We use induction on δ . We assume that $\delta \geq 6$ since the smaller cases are trivial. Consider the following linear program (without objective).

$$\begin{aligned} 0 &\leq x \leq 1 && \text{(box)} \\ \forall e \in E_T, \sum_{P: e \in P} x_P &\geq \ell_e := 3 && \text{(lb)} \\ \forall e \in E_T, \sum_{P: e \in P} x_P &\leq u_e := 3 + \text{degree}_H(e) - \delta(H). && \text{(ub)} \end{aligned}$$

This linear program has at least one feasible solution: $x = \frac{3}{\delta}$, which is easily seen to satisfy (box) and (lb), and which satisfies (ub):

$$3\text{degree}_H(e)/\delta \leq 3 + \text{degree}_H(e) - \delta \iff 0 \leq (\delta - 3)(\text{degree}_H(e) - \delta)/\delta.$$

We run the following iterated relaxation algorithm, where the initial LP is (box)–(ub).

1. Let x^* be an extreme point solution of the current LP.
2. If some path P has x_P^* equal to 0 or 1, then *fix* x_P : from now on we consider each appearance of x_P in the linear program to be a constant, equal to x_P^* .
3. If for some edge e there are only three non-fixed paths passing through e , then *relax* e : delete the constraints (lb) and (ub) for e .
4. Quit if all variables are fixed, otherwise return to the first step.

The LP remains feasible from iteration to iteration since fixing and relaxing both preserve feasibility of x^* . The main part of the proof will be to prove that the algorithm terminates. Assuming that this is so, here is how we complete the proof of Theorem 4.

Let x^f be the vector of final fixed values. We claim that every constraint (lb) and (ub) is almost satisfied by x^f , with at most an additive ± 2 violation. This is trivially true for each e that was never relaxed. If e is relaxed in some iteration, at that time $\sum_{P:e \in P} x_P^* \geq \ell_e$ holds, and the left-hand side is a sum of some fixed integers plus at most three non-fixed values strictly between 0 and 1; since ℓ_e is an integer, it follows that the fixed integers sum to at least $\ell_e - 2$. Hence these fixed values yield $\sum_{P:e \in P} x_P^f \geq \ell_e - 2$, no matter what happens to the non-fixed values over the remaining course of the algorithm. For (ub) a similar argument establishes that $\sum_{P:e \in P} x_P^f \leq u_e + 2$.

By our choice of $\ell_e = 3$, the inequalities $\sum_{P:e \in P} x_P^f \geq \ell_e - 2$ ensures that x^f is the indicator vector of a set cover \mathcal{E}_1 . Similarly, $\sum_{P:e \in P} x_P^f \leq u_e + 2$ ensures that $\mathbf{1} - x^f$ is the indicator vector of a family \mathcal{E}_2 of paths so that the hypergraph $H' = (E_T, \mathcal{E}_2)$ has $\delta(H') \geq \delta(H) - 5$. Apply Theorem 4 inductively to H' ; these $1 + \lfloor (\delta - 6)/5 \rfloor$ many set covers, together with \mathcal{E}_1 , give the desired result.

To prove Theorem 4 it remains only to prove that the algorithm terminates. In order to do this it is helpful to write the modified LP explicitly, collecting all fixed constants. Define ℓ'_e to be ℓ_e minus the x values of all fixed paths passing through e , and define u'_e similarly. Define the tree T' by contracting all relaxed edges, so $E_{T'}$ is the set of all non-relaxed edges. Let \mathcal{P}' be the non-fixed paths, updated to take the contraction into account. The modified LP, whose variables are x'_P for each non-fixed P , is

$$\mathbf{0} \leq x' \leq \mathbf{1} \text{ and } \forall e \in E_{T'}, \ell'_e \leq \sum_{P \in \mathcal{P}': e \in P} x'_P \leq u'_e. \quad (2)$$

Notice that extreme points of this LP correspond to extreme points of the one in the algorithm, so we denote both by x^* . If the algorithm does not fix any paths in a given iteration, notice that we have the strict inequality $\mathbf{0} < x^* < \mathbf{1}$.

Claim 15. *If x^* is an extreme point solution to (2) with $\mathbf{0} < x^* < \mathbf{1}$, then some $e \in E_{T'}$ lies on at most two paths from \mathcal{P}' .*

Proof. Since x^* is an extreme point solution and no box constraints are tight, there are disjoint sets $E_\ell, E_u \subseteq E_{T'}$ so that x^* is the unique solution to

$$\sum_{P:e \in P} x_P^* = \begin{cases} \ell'_e, & \forall e \in E_\ell; \\ u'_e, & \forall e \in E_u. \end{cases}$$

In particular, $|E_\ell| + |E_u| = |\mathcal{P}'|$ and the incidence matrix of $E_\ell \cup E_u$ versus \mathcal{P}' is invertible. From this point on, the argument is the same as that at the top of page 9 in the proof of [24, Lemma 4]. (The items $T = (V, E)$, E^* , D , c there correspond to T' , $E_\ell \cup E_u$, \mathcal{P}' , ℓ'/u' here. We make use of the integrality of ℓ' and u' , and linear independence in the aforementioned incidence matrix.) \square

Claim 15 proves that the algorithm either fixes or relaxes in each iteration, and therefore terminates. So the proof of Theorem 4 is complete. \square

This gives a 5-approximation algorithm for cd in this family of tree hypergraphs. From Holyer's result [21] mentioned earlier, it is NP-hard to approximate cd within a factor less than $3/2$ in the same family. We think Theorem 4 is not tight; the best upper bound on $\overline{\text{cd}}$ we know is $\lfloor (3\delta + 1)/4 \rfloor$.

For Theorem 4 the ground set is the edge set of an undirected tree, but it is equally natural to consider two other settings. One is to let the ground set be the vertex set of an undirected tree, with hyperedges corresponding to the vertices in paths. Another is to let the ground set be the arc set of a bidirected tree, with hyperedges corresponding to the arcs in directed paths. In the forthcoming full version of [24] it is shown that [24, Lemma 4] has analogues in these settings, with ± 6 in place of ± 2 , and from this we get that $\text{cd} \geq 1 + \lfloor (\delta - 1)/13 \rfloor$ in these settings.

For polychromatic numbers and systems of paths in trees, we have:

Theorem 5. $\overline{\text{p}}(\text{TREEEDGES}|\text{PATHS}, r) = \lceil r/2 \rceil$.

Proof. For the lower bound, colour the edges of the tree by giving all edges at level i the colour $i \bmod \lceil r/2 \rceil$. Since every length- r path in a tree has a monotonic subpath of length $\lceil r/2 \rceil$, we are done.

For the upper bound, by monotonicity of \bar{p} , it is enough to consider even $r = 2s$. Construct a spider tree that consists of a root vertex and $s + 2$ paths of s edges (each path has the root as one endpoint, and all other vertices are distinct). Define each of the $\binom{s+2}{2}$ leaf-leaf paths to be a hyperedge of the hypergraph; note that each one has size r . We must show that there is no polychromatic $(s + 1)$ -colouring. To see this, notice that each of the root-leaf paths must miss at least one colour, and by the pigeonhole principle two of them miss a common colour. The union of those two root-leaf paths is a hyperedge which is not coloured polychromatically, and we are done. \square

4 Small VC-Dimension

Recall the definition of VC-dimension: $S \subseteq V$ is *shattered* by a hypergraph if every $T \subseteq S$ can be obtained as an intersection of S with some hyperedge, i.e. if $\{S \cap E \mid E \in \mathcal{E}\} = 2^S$; and the VC-dimension of a hypergraph equals the size of the largest shattered set. The *dual VC-dimension* is the VC-dimension of the dual.

To show that primal and dual VC-dimension 2 is not enough to ensure cover-decomposability, it suffices to use the following construction of Pach, Tardos and Tóth [32]. Let T be a k -ary rooted tree, with k levels of vertices; so T has k^{k-1} leaves and $\sum_{i=0}^{k-1} k^i$ nodes in total. For each non-leaf node v we define its *sibling hyperedge* to be the k -set consisting of v 's children. For each leaf node v we define its *ancestor hyperedge* to be the k -set consisting of the nodes on the path from v to the root node. We let PTT_k denote the hypergraph consisting of all ancestor and sibling hyperedges. Pach et al. use a Ramsey-like argument to show $p(\text{PTT}_k) = 1$. Using this, we prove Theorem 7.

Theorem 7. *For the family of hypergraphs $\{H \mid \text{VC-dim}(H), \text{VC-dim}(H^*) \leq 2\}$, we have $\bar{p}(r) = 1$ for all r and $\overline{\text{cd}}(\delta) = 1$ for all δ .*

Proof. We will show that PTT_k has VC-dimension and dual VC-dimension at most 2. Thus PTT_k confirms the first part of Theorem 7 (since $k = r$) and PTT_k^* confirms the second part. A key observation is that in PTT_k ,

$$\text{any two distinct edges intersect in either 0 or 1 vertices.} \quad (3)$$

First, we bound the primal VC-dimension. Suppose for the sake of contradiction that there is a shattered vertex set $\{x, y, z\}$ of size 3. Since the set is shattered, there is a hyperedge E containing all of $\{x, y, z\}$. But by the definition of shattering, there must be another hyperedge $E' \neq E$ with $E' \cap \{x, y, z\} = \{x, y\}$. This contradicts (3), so we are done.

Second, we bound the dual VC-dimension. Suppose for the sake of contradiction that E, E', E'' are three hyperedges which are shattered in the dual. This implies that $E \cap E' \cap E''$ and $(E \cap E') \setminus E''$ are both nonempty. But this would imply $|E \cap E'| \geq 2$, contradicting (3). \square

In the remainder of this section we show that set systems with unit VC-dimension have large cover-decomposition and polychromatic numbers (Theorem 6). The following is a convenient way of looking at such hypergraphs.

Definition 16. *A hypergraph (V, \mathcal{E}) is called cross-free if the following holds for every pair $S, T \in \mathcal{E}$: at least one of $S \cap T, S \setminus T, T \setminus S$, or $V \setminus S \setminus T$ is empty.*

Observe that a hypergraph has VC-dimension 1 if and only if the dual hypergraph is cross-free. In the proofs of this section we will also use *laminar* hypergraphs, which are a subclass of cross-free hypergraphs.

Definition 17. *A hypergraph (V, \mathcal{E}) is called laminar if the following holds for every pair $S, T \in \mathcal{E}$: at least one of $S \cap T, S \setminus T$, or $T \setminus S$ is empty. (Equivalently, either $S \cap T = \emptyset, S \subseteq T$, or $T \subseteq S$.)*

We need the following fact; a one-line proof is that such hypergraphs are balanced, but an elementary proof is also an easy exercise.

Fact 18. *In a laminar hypergraph, $\text{cd} = \delta$.*

By duality, the following proposition proves the first half of Theorem 6.

Proposition 19. *For the family of cross-free hypergraphs, $\overline{\text{cd}}(\delta) = \lceil \delta/2 \rceil$.*

Proof. First, for each δ we demonstrate a cross-free hypergraph with minimum degree δ which cannot be decomposed into more than $\lceil \delta/2 \rceil$ covers. Take a hypergraph whose ground set V has $\delta+1$ elements, with one hyperedge $V \setminus \{v\}$ for each $v \in V$. Since each set cover has at least 2 hyperedges, we have $\overline{\text{cd}} \leq \lfloor (\delta+1)/2 \rfloor$, as needed.

To prove the other direction $\overline{\text{cd}} \geq \lfloor (\delta+1)/2 \rfloor$ for all cross-free hypergraphs, we use induction on δ . Clearly the bound holds for $\delta = 0$ or $\delta = 1$. For the inductive step, we have two cases. First, if the cross-free hypergraph (V, \mathcal{E}) has two sets $S, T \in \mathcal{E}$ for which $S \cup T = V$, then this is a cover with maximum degree 2 at each node. Thus we are done, since by induction $\overline{\text{cd}}(V, \mathcal{E} \setminus \{S, T\}) \geq \lceil (\delta-2)/2 \rceil$. Second, if the cross-free hypergraph has no such S, T , then in fact the cross-free hypergraph is laminar, and we are done by Fact 18. \square

Again by duality, the following proposition proves the second half of Theorem 6.

Proposition 20. *For the family of cross-free hypergraphs, $\overline{\text{p}}(r) = \lceil r/2 \rceil$.*

Proof. Consider again the hypergraph used in the first half of the proof of Proposition 19. Note that it is self-dual, e.g. since its incidence matrix is the all-ones matrix minus the identity matrix, which is symmetric. Thus, we deduce $\overline{\text{p}}(r) \leq \lfloor (r+1)/2 \rfloor$ for all r .

Next, we prove the other direction. We must show for all k , that if every hyperedge has size at least $2k-1$, then there is a polychromatic k -colouring. A key observation is the following: if \mathcal{E}^{\min} denotes the inclusion-minimal elements of \mathcal{E} , then a colouring is polychromatic for (V, \mathcal{E}) if and only if it is polychromatic for (V, \mathcal{E}^{\min}) . So we may reset $\mathcal{E} := \mathcal{E}^{\min}$, which does not affect r or cross-freeness. Moreover now \mathcal{E} is a clutter: there do not exist two different hyperedges $A, B \in \mathcal{E}$ for which $A \subset B$.

Lemma 21. *In a cross-free clutter (V, \mathcal{E}) , either all hyperedges are pairwise disjoint, or for every two hyperedges A, B we have $A \cup B = V$.*

Proof. From the definitions of “cross-free” and “clutter,” we first have that for every $S, T \in \mathcal{E}$, either $S \cap T = \emptyset$ or $S \cup T = V$. Note that it is impossible for three distinct sets $S, T, U \in \mathcal{E}$ to satisfy $S \cup T = V, S \cap U = \emptyset$ since this would imply $U \subset T$. Using this transitively, we obtain the lemma. \square

Using the lemma, Proposition 20 now boils down to two cases. The first case is that all hyperedges are pairwise disjoint. In this case a polychromatic k -colouring is easy to obtain: for each hyperedge E , just colour its $|E| \geq 2k-1 \geq k$ elements in any way that uses all k colours.

Finally, we deal with the case that for every two hyperedges A, B we have $A \cup B = V$. Rewriting, this means that \mathcal{E} is of the form $\{V \setminus S_1, V \setminus S_2, \dots, V \setminus S_m\}$ where $m \geq 2$, the S_i are pairwise disjoint, and $|V \setminus S_i| \geq 2k-1$ for all i . Here is a polychromatic colouring method which picks one colour class at a time. Repeat the following for $j = 1, \dots, k-1$: find any two vertices v, v' such that no S_i contains them both, and colour them with colour j . To see this is possible, note by induction that after iteration j , the number of uncoloured vertices in each $V \setminus S_i$ is at least $2k-1-2j$. After stage $k-1$, colour all remaining uncoloured vertices with colour k . \square

We remark that Theorem 6 may be viewed from the perspective of a family of tree-related hypergraphs, distinct from the ones we previously studied. Namely, Edmonds & Giles [13] showed that a hypergraph H is cross-free precisely when there exists a directed tree, such that each element of $V(H)$ occurs exactly once as a label on a node of the tree (a given node can have zero or multiple labels), and $\mathcal{E}(H)$ consists of those label sets on the half-tree “pointed to” by each directed edge of the tree.

5 $O(R)$ -approximation for Sensor Cover

The *sensor cover* problem is a generalization of cover-decomposition. The input is a hypergraph plus, for each hyperedge E , a *duration* ℓ_E ; a *schedule* is an activation time a_E for each hyperedge, and it has *coverage* T if for all $0 \leq t \leq T$ and all $v \in V$, some edge E has $v \in E$ and $t \in [a_E, a_E + \ell_E]$. We are interested in finding a schedule with maximum coverage. Note that in the special case where all durations are unit, the maximum coverage equals the cover-decomposition number.

Let $\bar{\delta} = \min_v \sum_{E \ni v} \ell_E$, the minimum duration-weighted degree. Clearly the maximum schedule coverage is at most $\bar{\delta}$. The most prominent results for this setting are from [18] (the conference version): it is shown that every interval hypergraph has a schedule of coverage $\Omega(\bar{\delta})$; using this as a subroutine it is shown that for hypergraphs corresponding to translates of any fixed open convex polygon covering points in \mathbb{R}^2 , a schedule of $\Omega(\bar{\delta})$ always exists.

Let R still mean the maximum number of vertices in any hyperedge, independent of the durations. We then obtain the following result.

Theorem 22. *Every sensor cover instance admits a schedule of coverage $\Omega(\bar{\delta}/R)$.*

Proof. Without loss of generality, scale all durations uniformly so that $\bar{\delta} = 1$. Let α be a parameter to be fixed later in the range $O(R)$, so that we will seek a schedule of coverage $1/\alpha$.

Clearly if there are any hyperedges of duration at least $1/\alpha$ we can schedule them at time zero, satisfying all their contained vertices; so without loss of generality all durations are $< 1/\alpha$. Then, round down durations of the remaining hyperedges to the closest member of $\{2^{-1}/\alpha, 2^{-2}/\alpha, \dots\}$. Let level i consist of those hyperedges whose new durations are $2^{-i}/\alpha$, and let $d_i(v)$ be the degree of vertex v in level i . Note that after rounding, the duration-weighted degree of each vertex is at least $1/2$, so

$$\sum_{i \geq 1} \frac{d_i(v)}{2^i \alpha} \geq 1/2. \quad (4)$$

In each level we will use:

Claim 23. *In a hypergraph, we can assign each hyperedge to a vertex it contains, so that for each v , the number of hyperedges assigned to it is at least $\lfloor d(v)/R \rfloor$.*

Proof. The naive LP formulation for this assignment problem (with integral data) is well-known to be integral. There is a feasible fractional assignment (fractionally assign each E an amount $1/|E|$ to each vertex it contains), so there is also a feasible integral assignment. \square

In each level, find the assignment specified by Claim 23. We will be done if we can pick $\alpha = O(R)$ so that, adding up all levels, the total duration assigned to each vertex is at least $1/\alpha$: a satisfactory schedule can then be obtained by, for each vertex, scheduling its assigned edges one after the other. Since $\lfloor x/R \rfloor \geq (x - (R - 1))/R$ for integral x and R , the duration of the edges assigned to each v is at least

$$\begin{aligned} \sum_{i \geq 1} \frac{\lfloor d_i(v)/R \rfloor}{2^i \alpha} &\geq \sum_{i \geq 1} \frac{(d_i(v) - (R - 1))/R}{2^i \alpha} = \frac{1}{R} \sum_{i \geq 1} \frac{d_i(v)}{2^i \alpha} - \frac{R - 1}{R} \sum_{i \geq 1} \frac{1}{2^i \alpha} \\ &\geq \frac{1}{R} \frac{1}{2} - \frac{R - 1}{R} \frac{1}{\alpha}, \end{aligned}$$

where in the last inequality we used (4) and the geometric series sum formula. We need precisely that the right-hand side above is at least $1/\alpha$. A short calculation reveals that $\alpha = 4R - 2$ will do, and completes the proof. \square

We remark that, by tweaking the scaling factor of 2 in the proof above, the final ratio α can be slightly improved from $4R - 2$ to $2R - 1 + 2\sqrt{R^2 - R}$.

6 Open Problems

The major open problem in our study is whether the sensor cover result Theorem 22 can be improved to a schedule of coverage $\Omega(\bar{\delta}/\ln R)$, in line with the tight results from Section 2. We also do not know if in TREEEDGES|PATHS, a schedule of coverage $\Omega(\bar{\delta})$ is always possible. For the non-scheduling version, it would be interesting to know more about $\overline{\text{cd}}(\text{HYP}(R), \delta)$ in the regime $\delta = \Theta(\ln R)$, or for small values of δ and R .

Two outstanding open geometric problems that predate our work are to determine whether $\bar{\text{p}}(\delta) = \Omega(\delta)$ for the family of unit discs, and whether the family of all axis-aligned squares is cover-decomposable.

Pálvölgyi [34] poses two nice combinatorial questions: first, is there a function f so that in hypergraph families closed under edge deletion and duplication, $\overline{\text{cd}}(\delta_0) \geq 2$ implies $\overline{\text{cd}}(f(\delta_0)) \geq 3$? This compelling question, which would give a unified explanation for much of the existing literature, is open even for $\delta_0 = 2$; no counterexamples are known to the hypothesis $f(\delta_0) = O(\delta_0)$, i.e. that cover-decomposability implies p grows linearly with δ . Second, consider an $m \times r$ matrix of integers which is weakly increasing from left to right and top to bottom. If we think of each row as a set of size r , giving a hyperedge on ground set \mathbb{Z} , do these *shift chains* have Property B for large enough r ? It is known [34] that $r = 3$ is not enough.

Acknowledgments

We wish to thank Jessica McDonald, Dömötör Pálvölgyi, and Oliver Schaud for helpful discussions on these topics. We thank the ESA and SIDMA referees for their suggestions, in particular for the nicer proof of the second half of Theorem 5.

References

- [1] NOGA ALON, ROBERT BERKE, KEVIN BUCHIN, MAIKE BUCHIN, PÉTER CSORBA, SASWATA SHANNIGRAHI, BETTINA SPECKMANN, AND PHILIPP ZUMSTEIN, *Polychromatic colorings of plane graphs*, Discrete & Computational Geometry, 42 (2009), pp. 421–442. Preliminary version appeared in *Proc. 24th SOCG*, pages 338–345, 2008.
- [2] NOGA ALON AND JOEL H. SPENCER, *The Probabilistic Method (3rd Edition)*, Wiley, New York, 2008.
- [3] GREG ALOUPIS, JEAN CARDINAL, SÉBASTIEN COLLETTE, SHINJI IMAHORI, MATIAS KORMAN, STEFAN LANGERMAN, ODED SCHWARTZ, SHAKHAR SMORODINSKY, AND PEROUZ TASLAKIAN, *Colorful strips*, Graphs and Combinatorics, 27 (2011), pp. 327–339. Preliminary version in *Proc. 9th LATIN*, pages 2–13, 2010.
- [4] GREG ALOUPIS, JEAN CARDINAL, SÉBASTIEN COLLETTE, STEFAN LANGERMAN, DAVID ORDEN, AND PEDRO RAMOS, *Decomposition of multiple coverings into more parts*, Discrete Comput. Geom., 44 (2010), pp. 706–723. Preliminary version in *Proc. 20th SODA*, pages 302–310, 2009.
- [5] LARS DØVLING ANDERSEN, *Lower bounds on the cover-index of a graph*, Discrete Mathematics, 25 (1979), pp. 199–210.
- [6] PAUL BALISTER, BÉLA BOLLOBÁS, AMITES SARKAR, AND MARK WALTERS, *Sentry selection in wireless networks*, Adv. in Appl. Probab., 42 (2010), pp. 1–25.
- [7] JÓZSEF BECK AND TIBOR FIALA, *“Integer-making” theorems*, Discrete Applied Mathematics, 3 (1981), pp. 1–8.
- [8] ERIC BLAIS, AMIT WEINSTEIN, AND YUICHI YOSHIDA, *Semi-strong coloring of intersecting hypergraphs*. arXiv:1203.2868, 2012.
- [9] BÉLA BOLLOBÁS, DAVID PRITCHARD, THOMAS ROTHVOSS, AND ALEX D. SCOTT, *Cover-decomposition and polychromatic numbers*, in *Proc. 19th ESA*, 2011, pp. 799–810.

- [10] ADAM L. BUCHSBAUM, ALON EFRAT, SHAILI JAIN, SURESH VENKATASUBRAMANIAN, AND KE YI, *Restricted strip covering and the sensor cover problem*, in Proc. 18th SODA, 2007, pp. 1056–1063.
- [11] TIMOTHY M. CHAN, ELYOT GRANT, JOCHEN KÖNEMANN, AND MALCOLM SHARPE, *Weighted capacitated, priority, and geometric set cover via improved quasi-uniform sampling*, in Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '12, SIAM, 2012, pp. 1576–1585.
- [12] GRUIA CĂLINESCU, CHANDRA CHEKURI, AND JAN VONDRÁK, *Disjoint bases in a polymatroid*, Random Struct. Algorithms, 35 (2009), pp. 418–430.
- [13] JACK EDMONDS AND RICK GILES, *A min-max relation for submodular functions on graphs*, in Studies in Integer Programming (1975 Bonn, Germany), vol. 1 of Annals of Discrete Mathematics, North-Holland, 1977, pp. 185–204.
- [14] MÁRTON ELEKES, TAMÁS MÁTRAI, AND LAJOS SOUKUP, *On splitting infinite-fold covers*, Fund. Math., 212 (2011), pp. 95–127.
- [15] PAUL ERDŐS AND LÁSZLÓ LOVÁSZ, *Problems and results on 3-chromatic hypergraphs and some related questions*, in Infinite and Finite Sets (Coll. Keszthely, 1973), A. Hajnal, R. Rado, and V. T. Sós, eds., vol. 10 of Colloq. Math. Soc. János Bolyai, North-Holland, 1975, pp. 609–627.
- [16] URIEL FEIGE, MAGNÚS M. HALLDÓRSSON, GUY KORTSARZ, AND ARAVIND SRINIVASAN, *Approximating the domatic number*, SIAM J. Comput., 32 (2002), pp. 172–195. Preliminary version appeared in Proc. 32nd STOC, pages 134–143, 2000.
- [17] RADOSLAV FULEK, *Coloring geometric hypergraph defined by an arrangement of half-planes*, in Proc. 22nd CCCG, 2010, pp. 71–74.
- [18] MATT GIBSON AND KASTURI VARADARAJAN, *Optimally decomposing coverings with translates of a convex polygon*, Discrete & Computational Geometry, 46 (2011), pp. 313–333. 10.1007/s00454-011-9353-9.
- [19] RAM PRAKASH GUPTA, *On decompositions of a multi-graph into spanning subgraphs*, Bull. Amer. Math. Soc., 80 (1974), pp. 500–502.
- [20] RAM P. GUPTA, *On the chromatic index and the cover index of a multigraph*, in Theory and Applications of Graphs: Int. Conf. Kalamazoo, May 11–15, 1976, Yousef Alavi and Don R. Lick, eds., vol. 642 of Lecture Notes in Mathematics, Springer Verlag, 1978, pp. 204–215.
- [21] IAN HOLYER, *The NP-completeness of edge-coloring*, SIAM J. Comput., 10 (1981), pp. 718–720.
- [22] BALÁZS KESZEGH, *Weak conflict free colorings of point sets and simple regions*, in Proc. 19th CCCG, 2007. Extended version “Coloring half-planes and bottomless rectangles” at arXiv:1105.0169.
- [23] BALÁZS KESZEGH AND DÖMÖTÖR PÁLVÖLGYI, *Octants are cover-decomposable*, Discrete & Computational Geometry, 47 (2012), pp. 598–609. Preliminary version in Proc. 6th EuroComb, pages 499–504, 2011.
- [24] JOCHEN KÖNEMANN, OJAS PAREKH, AND DAVID PRITCHARD, *Max-weight integral multicommodity flow in spiders and high-capacity trees*, in Proc. 6th WAOA, 2008, pp. 1–14. Extended version “Multicommodity Flow in Trees: Packing via Covering and Iterated Relaxation” submitted to a journal, 2011+.
- [25] LÁSZLÓ LOVÁSZ AND MICHAEL D. PLUMMER, *Matching Theory*, Akadémiai Kiadó, Budapest, 1986.
- [26] PETER MANI-LEVITSKA AND JÁNOS PACH, *Decomposition problems for multiple coverings with unit balls*. Manuscript, 1987.

- [27] ———, *Unsplittable multiple coverings with balls*. Manuscript, available online, 1988.
- [28] ROBIN A. MOSER AND GÁBOR TARDOS, *A constructive proof of the general Lovász Local Lemma*, J. ACM, 57 (2010). Preliminary version in *Proc. 41st STOC*, 343–350, 2009.
- [29] JÁNOS PACH, *Decomposition of multiple packing and covering*, in Kolloquium über Diskrete Geometrie, Salzburg, 1980, Inst. Math. U. Salzburg, pp. 169–178.
- [30] JÁNOS PACH, *Covering the plane with convex polygons*, Discrete and Computational Geometry, 1 (1986), pp. 73–81.
- [31] JÁNOS PACH AND GÁBOR TARDOS, *Tight lower bounds for the size of epsilon-nets*, in Proc 27th SoCG, 2011, pp. 458–463. arXiv:1012.1240.
- [32] JÁNOS PACH, GÁBOR TARDOS, AND GÉZA TÓTH, *Indecomposable coverings*, Canadian Mathematical Bulletin, 52 (2009), pp. 451–463. Preliminary version in *Proc. 7th CJDGCGT* (2005), pages 135–148, 2007.
- [33] JÁNOS PACH AND GÉZA TÓTH, *Decomposition of multiple coverings into many parts*, Comput. Geom. Theory Appl., 42 (2009), pp. 127–133. Preliminary version appeared in Proc. 23rd SoCG, pages 133–137, 2007.
- [34] DÖMÖTÖR PÁLVÖLGYI, *Decomposition of Geometric Set Systems and Graphs*, PhD thesis, École Polytechnique Fédérale de Lausanne, 2010. arXiv:1009.4641.
- [35] ———, *Indecomposable coverings with concave polygons*, Discrete and Computational Geometry, 44 (2010), pp. 577–588. 10.1007/s00454-009-9194-y.
- [36] DÖMÖTÖR PÁLVÖLGYI AND GÉZA TÓTH, *Convex polygons are cover-decomposable*, Discrete and Computational Geometry, 43 (2010), pp. 483–496.
- [37] ALEXANDER SCHRIJVER, *Supermodular Colourings*, in Matroid Theory, L. Lovász and A. Recski, eds., North-Holland, 1985, pp. 327–343.
- [38] ALEXANDER SCHRIJVER, *Combinatorial optimization*, Springer, New York, 2003.
- [39] SHAKHAR SMORODINSKY AND YELENA YUDITSKY, *Polychromatic coloring for half-planes*, J. Comb. Theory, Ser. A, 119 (2012), pp. 146–154. Preliminary version in *Proc. 12th SWAT*, pages 118–126, 2010.
- [40] GÁBOR TARDOS AND GÉZA TÓTH, *Multiple coverings of the plane with triangles*, Discrete and Computational Geometry, 38 (2007), pp. 443–450.
- [41] CARSTEN THOMASSEN, *The even cycle problem for directed graphs*, J. Amer. Math. Soc., 5 (1992), pp. 217–229.